

PCT

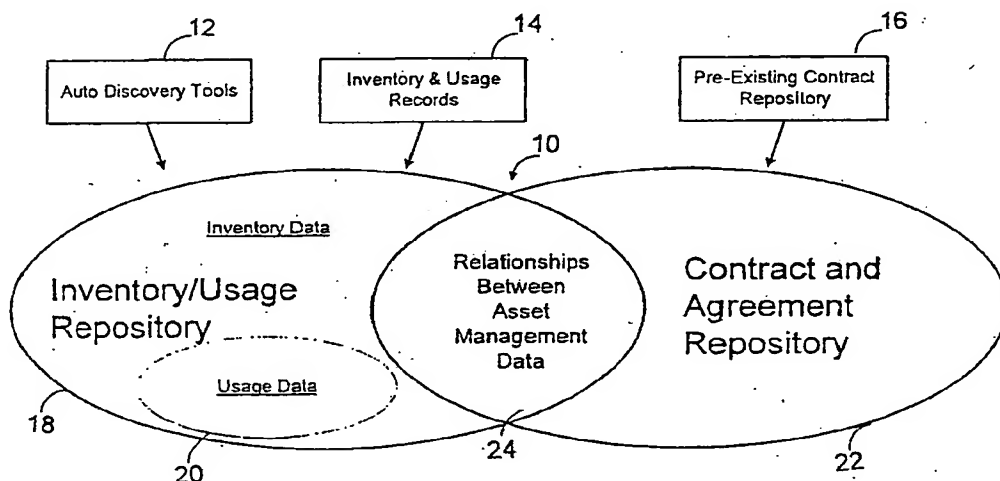
WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁷ : G06F 1/00	A1	(11) International Publication Number: WO 00/52559 (43) International Publication Date: 8 September 2000 (08.09.00)
(21) International Application Number: PCT/US00/05664 (22) International Filing Date: 3 March 2000 (03.03.00) (30) Priority Data: 60/123,531 5 March 1999 (05.03.99) US Not furnished 2 March 2000 (02.03.00) US (71) Applicant: ISOGON CORPORATION [US/US]; 330 Seventh Avenue, New York, NY 10001 (US). (72) Inventors: BARRITZ, Robert; Isogon Corporation, 330 Seventh Avenue, New York, NY 10001 (US). KASSAN, Peter; Isogon Corporation, 330 Seventh Avenue, New York, NY 10001 (US). HELLBERG, Per; Isogon Corporation, 330 Seventh Avenue, New York, NY 10001 (US). (74) Agents: WEINER, Samuel, H. et al.; Ostrolenk, Faber, Gerb & Soffen, LLP, 1180 Avenue of the Americas, New York, NY 10036 (US).		(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG). Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>

(54) Title: **METHOD AND PROCESS FOR DISPLAYING SOFTWARE PRODUCT INVENTORY AND USAGE INFORMATION CORRELATED WITH LICENSE AGREEMENT INFORMATION**



(57) Abstract

A software tool and database that enables a user to view information, on the one hand, pertaining to hardware and software product inventory and usage, and, on the other hand, to license agreement information in such way that the two sets of information are correlated and associated. The tool presents inventory and usage information based on agreement criteria and/or agreement information based on inventory and usage criteria in an integrated fashion that permits the user to switch from one kind of information to the other kind in a variety of convenient ways. The main functional blocks of the software tool include a data acquisition tool, a reconciliation and association tool, and a processing/display/output tool.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon	KR	Republic of Korea	PL	Poland		
CN	China	KZ	Kazakhstan	PT	Portugal		
CU	Cuba	LC	Saint Lucia	RO	Romania		
CZ	Czech Republic	LI	Liechtenstein	RU	Russian Federation		
DE	Germany	LK	Sri Lanka	SD	Sudan		
DK	Denmark	LR	Liberia	SE	Sweden		
EE	Estonia			SG	Singapore		

METHOD AND PROCESS FOR DISPLAYING SOFTWARE
PRODUCT INVENTORY AND USAGE INFORMATION
CORRELATED WITH LICENSE AGREEMENT INFORMATION

BACKGROUND OF THE INVENTION

5 The present invention relates to Software and Hardware Inventory and Usage and more particularly to a method and process for displaying software product inventory and usage information correlated with license agreement information.

10 The invention builds upon the disclosure in the present assignees U.S. Patent No. 5,499,340, the contents of which are incorporated by reference herein.

 A variety of techniques are available for gathering and reporting on the inventory and usage of software products and the hardware (computers) that they are
15 installed and execute upon.

 In some environments (e.g., PC networks), software tools are available that conduct a survey to discover the computers and related hardware devices. Some of these
20 same tools, as well as other products that do not do such a hardware survey, also survey the executable software programs ("modules," "executable files") on a computer or a network of computers and then automatically identify (on the basis of a knowledge base of file names,
25 identifying strings, "footprints," etc.) which software products they comprise.

 Yet other tools allow users to specify products that users are installing or have installed, implicitly or

- 2 -

explicitly indicating which programs they comprise. In addition, some of these tools identify the particular versions and/or releases found as well as what product suites each one belongs to. Such inventory information is often enhanced to include identification of the particular processor on which each product is installed, its physical location (by country, region, state, county, city, building, etc.) and organizational subdivisions (such as company, division, and department) with which the processor, the product, or the particular installation of the product is associated. Thereafter, such inventory information can be treated on a consolidated basis, with the ability to select and manipulate on the basis of such enhanced identification.

Using such a tool (or other such tools), the usage or execution of such software products (which have been automatically identified or manually specified by the user upon installation or at a later time) can be monitored. A variety of techniques are available to monitor the usage of other programs, including: intercepting the operating system component that manages or initiates program execution; making use of interfaces provided by the operating system for such purposes; "wrapping" or "prefacing" program modules with code to invoke or perform the monitoring functions; intercepting operating system logging activity; reading operating system logs after the fact; and other such techniques. A number of other tools are available that monitor such module usage without themselves relating modules to product identity.

- 3 -

Thereafter, a variety of product inventory and usage information can be displayed, printed, exported, etc. For example:

Installed products by processor and library

5 Installed products by vendor

Libraries by installed product, upon which each product is installed,

Product usage by product

Product usage by product and library

10 Product usage by user and product

Product usage by user, product, processor, and library

License Agreements

Other software tools exist that allow users to enter
15 into a database information pertaining to contractual license agreements ("contracts," "agreements") for the user's software products.

This provides a convenient place to record such information as vendor or distributor, contact
20 information, software and/or hardware products covered by the agreement, computers on which the software is authorized to run, number of authorized users, maintenance terms, renewal date, budgets and forecasts, invoice and payment history, key terms and conditions,
25 links to electronic versions of documents, etc.

As license agreements are often interrelated in complex ways, such tools typically allow the user to organize and identify these agreements in terms of master agreements, sub-agreements, addenda, etc. Some such
30 products also provide: automated "alerts" (signals that

an important event is occurring); financial capabilities such as chargebacks, links to accounting systems, etc.; information about vendor Internet address (allowing the user to automatically open a browser and display information from the vendor's Web site); and other capabilities.

Existing software tools which offer asset management capabilities include ARGIS from Janus Technologies, The Contract Tickler from ARM Group and ContractManagement from Global Business Center. These existing products are not believed to provide linked associations of data from pre-existing repositories and many-to-many linkages. Nor do they provide reconciliation of data from multiple repositories in an automatic or substantially automatic fashion.

The Problem

Although both kinds of information--inventory and usage information on the one hand and agreement information on the other--are useful by themselves, they would be more meaningful and useful if they were integrated, related, coordinated, and reconciled.

For example, when the time comes to decide whether to renew a maintenance contract for a particular product from a particular vendor, it is highly desirable to be able to look at installation and usage information about that product. As another example, when looking at the usage patterns for a particular product, it is highly desirable to be able to look at the license agreement or agreements that pertain to it, for opportunities to save money by canceling a product or consolidating its usage

onto fewer processors.

However, at present, using existing tools, it is awkward to relate the two types of information. The user must look up the contract in the license agreement tool
5 and then manually look up the corresponding inventory and usage information in the other tool (or vice versa).

With a single installation having hundreds of agreements, thousands of software products, and possibly tens of thousands of computers and users, this process
10 can be laborious, repetitive, error-prone, expensive, and impractical. Time is often critical—if the information cannot be obtained and evaluated before a particular date (e.g., a deadline for canceling or renegotiating a contract), it is essentially useless.

15 The inability to cope with an unlimited quantity of such data in a reliable, timely and cost-effective way makes it impractical to perform such analysis and correlation for all vendors and all products on a regular basis — only those few occasions where the costs are high.
20 enough and the number of data items small enough to make it cost-justifiable.

Another problem in going back and forth between two unintegrated tools is the difficulty arising when reconciling items (such as a product, vendor, processor,
25 user, etc.) named or treated differently in the two tools. Although the user can, of course, make a decision, for example, as to which product in one tool corresponds to which product in the other, each user must make such a decision on a case-by-case basis, and there
30 is no convenient way to create an automatic linkage. This

- 6 -

means that, for example, the next time the user pulls up the same information pertaining to a particular license agreement for a particular product, he or she will have to repeat the same process. Furthermore, other users of
5 the same tools will each individually have to repeat the same process, possibly making different decisions.

Some tools may provide the ability to create multiple views ("windows"). The operating system itself may provide the capability to simultaneously use two or
10 more unrelated programs simultaneously in multiple views ("windows"), which allows the user to use multiple tools, an inventory/usage tool in one window and a license agreement tool in another window. The user would be able to relate the information in one tool to the information
15 in the other visually, and, perhaps even to cut information from one and copy it into the other. However, no tool allows the user to correlate the metrics of a license agreement with the actual (or proposed) operation of the software itself. Users would benefit if they were
20 able to conveniently answer questions such as the following:

For each product found in inventory which agreement governs its use?

For all products found in inventory
25 which products have an agreement associated with them?
which products have no agreement associated with them?

For all products meeting user-specified criteria
30 (e.g., found on a particular processor, from a particular

vendor, etc.), which agreements govern their use?

For all products with usage above or below a user-specified level over a user-specified time (e.g., unused for the past year) and meeting user-specified criteria
5 (e.g., found on a particular processor, from a particular vendor, etc.), what agreements govern their use?

For all products whose maintenance contracts are within a user-specified time of cancellation or renewal and meeting user-specified criteria (e.g., found on a
10 particular processor, from a particular vendor, etc.), what are the recent and current usage levels?

For each product for which the user has an agreement and meeting user-specified criteria (e.g., found on a particular processor, from a particular vendor, etc.), on
15 which processors is the product installed and what are the recent and current usage levels?

For all products for which the user has an agreement and meeting user-specified criteria (e.g., from a particular vendor, etc.), which products have not been
20 found in inventory?

For each product having a license agreement specifying a maximum number of MIPS (millions of instructions per second, a measurement of relative processing power) and meeting user-specified criteria
25 (e.g., found on a particular processor, from a particular vendor, etc.), what is the total MIPS of the processors for which the product is installed? As used herein, MIPS is an inclusive term which also refers to MSU (million service units, CPU speed, drystone, whetstone, or any
30 other measure of computing power.)

- 8. -

For each product for which the user has an agreement that specifies a particular number of processors and meeting user-specified criteria (e.g., from a particular vendor, etc.):

5 What is the number of processors on which it is installed, and is that number higher or lower than the agreement authorizes?

 What is the number of processors on which it has been used during a user-specified period of time, and is
10 that number higher or lower than the agreement authorizes?

 For each product for which the user has an agreement that specifies that the product can be installed on one or more particular processors and meeting user-specified
15 criteria (e.g., from a particular vendor, etc.):

 On what processors is the product installed, and is any such use not authorized by the agreement?

 On what processors has the product been used during a user-specified period of time, and is any such use not
20 authorized by the agreement?

 Likewise, the user would be able to be selective in his queries by optionally selecting processors, specified locations, and/or within user-specified organizational subdivisions to narrow the desired results.

25 SUMMARY OF THE PRESENT INVENTION

 The present invention is designed to make it easy to answer the above and other such questions, and to overcome the problems inherent in not having the information pertaining to hardware and software product

- 9 -

inventory and usage (on the one hand) and agreement information (on the other hand) coordinated.

The present invention is a software program and/or a knowledge base, hereinafter, the "tool," that enables the user to view information on the one hand pertaining to hardware and software product inventory and usage, and, on the other hand, license agreement information in such a way that the two sets of information are correlated and associated. The tool presents inventory and usage information based on agreement criteria and/or agreement information based on inventory and usage criteria in an integrated fashion that permits the user to switch from one kind of information to the other kind in a variety of convenient ways. The key components of the tool are the main software blocks which carry out the following main functional tasks: Data acquisition; Reconciliation and association; and Processing, display, and output.

As used herein, the term computer or computers -used by itself- means all types of computers from single to multiple PCs to midrange computers, UNIX, mainframe, networked, etc.

Other features and advantages of the present invention will become apparent from the following description of the invention which refers to the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 illustrates a relationship between agreements, inventory and usage data.

Figure 2 is a software screen illustrating assets by

- 10 -

contract comparisons.

Figure 3 is another software screen providing an asset usage by contract view.

5 Figure 4 is a conceptual data model of the present invention.

Figure 5 illustrates a table of associations.

Figure 6 is another software screen illustrating asset usage by contract and location.

DETAILED DESCRIPTION OF THE INVENTION

10 Introduction

The present invention is a software program, hereinafter, the "tool," that enables the user to view information on the one hand pertaining to hardware and software product inventory and usage, and, on the other
15 hand, license agreement information in such a way that the two sets of information are correlated and associated. The tool presents inventory and usage information based on agreement criteria and/or agreement information based on inventory and usage criteria in an
20 integrated fashion that permits the user to switch from one kind of information to the other kind in a variety of convenient ways.

The tool can be understood in terms of its major functional areas:

25 Data Acquisition
Reconciliation and Association
Processing, Display, and Output.

In the present invention, two or more sets of data (inventory/usage and agreement information) must be made

- 11 -

known to the tool. There are several ways that this can be done, such as:

The tool can itself provide the functionality already described to obtain hardware, software inventory
5 and/or usage information.

The information can be imported from one or more other products that provide that functionality (via reading a file or a database).

10 The tool can share a database with one or more products that already provide that functionality.

The tool can provide a facility by which the user can manually enter the information. (This is particularly applicable to the agreement information, as well as such information as the number of MIPS of each processor or
15 processor subdivision.)

The tool can make use of APIs (Application Program Interfaces) provided by other products to obtain agreement information or inventory and usage information from the databases or data files of those other products.

20 The tool can provide APIs for other products to use to supply agreement information or inventory and usage information to the present program whenever such information is created or obtained by those other products.

25 With reference to Figure 1, the present invention may interface with and/or acquire the bulk of its data from existing software inventory/usage and contract/agreement collectors as well as from data entered manually by operators. In one aspect thereof,
30 the invention comprises a specially constructed database

- 12 -

10 which produces relationships between diverse pieces of
asset management data. The database 10 contains
information including inventory data 18 and usage data
20. This inventory and usage data can be obtained from
5 auto discovery tools 12 and inventory and usage records
14, as shown. The database 10 further contains a
contract and agreement repository 22, the information for
which can be obtained from a pre-existing contract
repository 16. The aim of the invention is to create
10 data sets defining relationships between asset management
data, such as between inventory data and contract and
agreement data, as shown by the overlapped region 24 in
Figure 1, which represents data in the different data
sets that are interrelated.

15 The size of the relationships data 24 relative to
the size and content of the inventory data/usage data 18
or to the contract and agreement data 22 is dependent on
the ability of the software to identify those
relationships, as well as on the existence, in the first
20 instance, of such relationships. For example, at some
installations much of the inventory data may describe
home-grown software products which are not covered by
contract and agreement data. The usefulness of the
present invention becomes more pronounced as the overlap
25 region 24 increases in size and is quite significant when
the linked data records in the region 24 constitute at
least 25% of the number of records in either of the
databases 18 or 20.

30 With further reference to the drawings, each set of
data can be considered a repository of such data that is,

- 13 -

presumably, created, modified and updated by processes external to the present invention. These repositories, while conceptually distinct, may be stored in separate databases or a single database.

5 The contract/agreement repository (CDB) 22 is a database of contractual information organized according to a unique agreement identification code (ID), and further according to the type of data it contains, e.g., the asset, asset group, invoice, purchase order and
10 contract. Asset is defined to include but not be limited to hardware, software products, maintenance and service agreements. An asset group is a user-selected set of assets.

15 Each record or group of records contained in the CDB provides details of an agreement (lease, contract, license, etc.). Typically, such records include the following information:

- Time and date of transaction
- Vendor
- 20 • Name, address, phone number, email address, etc. of vendor
- General description of the asset (computer, DASD, software, maintenance, etc.)
- Manufacturer, model, year, options, etc.
- 25 • Warranties or support purchased
- Contract terms and conditions (purchase, lease, rental, etc.)
- Contact information (support personnel, phone numbers, email addresses, etc.)

30 These fields may also include more detailed

information such as: licensing fees; number of persons authorized, MIPS authorized, number of invocations authorized, annual maintenance fees, and product/sub-product identifications.

5 The primary purpose of the CBD is to maintain a complete database of contract information. If a data element is not already supported by the CDB, the present invention will, optionally, create and maintain the data for that element. For example, if one vendor's CDB does
10 not support asset groups, the present invention will provide the user with the capability of defining and maintaining asset groups.

 In addition to electronic storage of the contract information, the present invention also provides the user
15 with one or more ways to further store and/or reference an agreement:

1. An image of the original document is stored and/or referenced (i.e., file location, hyperlink, etc.). The image may have originated
20 from a facsimile, optical scanner, digital photo, or image file generated by the user and/or vendor. Storage formats may be bitmap, JPEG, TIFF, or other known storage format.
2. Optical character recognition (OCR) techniques
25 may be applied to a scanned document and the resulting data stored.
3. Electronic copy of the document. For example, a word processing file, spreadsheet, XML file, etc.
- 30 4. A description of the physical location of the

- 15 -

original hardcopy document.

Figure 2 is a software screen 26 providing an example of asset information viewed by contract.

Reconciliation and Association

5 After the two types of information have been acquired, the next crucial step is the means to associate the two.

Product identification based on inventory is often presented in different terms than the identification
10 based on agreements. For example, product or vendor names may not be identical; what is considered a single product from one point of view may be considered to consist of two or more products from the other; there may be other discrepancies and mismatches; etc.

15 Furthermore, even after all such inconsistencies, discrepancies, and mismatches have been resolved, there may be products for which there is license agreement information but no inventory information (e.g., the product has not yet been installed; it has been removed;
20 or the licensing information is in error). Similarly, there may be inventory information but no license agreement information (e.g., the user doesn't have an agreement for the product; the product doesn't require an agreement; or the information is missing). Such products
25 and agreements need to be identified, the reasons for such discrepancies need to be resolved, and the appropriate records must be created or deleted.

The present invention includes facilities by which all these tasks can be accomplished in a straightforward,
30 substantially automated manner. By substantial

- 16 -

automation is meant a process which reduces the time to produce data correlation/reconciliation to below about one fifth and preferably one tenth the time that the task would consume if it were carried out manually. The following list is meant to be illustrative and not exhaustive. Some of these facilities can be considered optional.

A facility for concurrently presenting two or more lists of items (e.g., products, vendors, processors, users, etc.) sorted by product name, product identification number, vendor name, vendor identification, or other identifying feature in a display.

Optionally, a facility by which the tool automatically establishes an association between items in one list (e.g., products, vendors, processors, users, etc.) with items in another list of such items whenever product name or product identification number match. These would be subject to user confirmation either globally or item-by-item.

Optionally, a facility by which the tool determines possible associations or associations that are most probable. This would be based upon the heuristics of evaluating the similarity of item name, identification number, or other such criteria between items in one list and items in another list. [Techniques that could be applied include: Soundex coding, pattern matching, dictionary lookup, and/or Fuzzy Logic.] The results would be displayed for user selection, confirmation, or rejection. For each item to be associated, the facility

- 17 -

provides as many candidates as are appropriate, ranked alphabetically or in likelihood of match. All such candidate associations can be displayed in a distinctive color; the user could tab or scroll directly from one such candidate association to the next and confirm or reject each one in turn, or could highlight and accept or reject several such candidate associations at once

A facility by which the user specifies an association between one or more items in one list with an item in another list. The user might use a pointing device such as a computer mouse to highlight one or more items on the one list of products and one item on the other list of products. Optionally, the tool would present an opportunity for the user to confirm the association before it was processed.

A facility for making changes to the associations already established. There are several ways in which these associations can be implemented.

One method is for each record in a table to be linked with a corresponding record in another table. Into that field, for each record, is written the index (pointer) of the corresponding record (that is, the record to be linked to) in the other table. This can be done reciprocally, so that each table has a field that points to records in the other table, as illustrated below:

LICENSE AGREEMENT PRODUCTS TABLE				
Record Number	Product Name	Vendor Name	Effective Date	Inventory/Usage Record Number
0000001	OS/390	IBM	05/04/89	0000079

0000002	CICS	IBM	05/04/89	0000063
0000003	ABEND-AID	CA	01/01/93	0000042
...				
...(etc.)				

- 19 -

INVENTORY/USAGE PRODUCTS TABLE			
Record Number	Product Name	Vendor Name	License Agreement Record Number
....			
0000042	CA-ABEND Aid	Computer Associates Int'l.	0000003
...			
0000063	Cust. Info. Control Sys.	International Business Machines Corporation	0000002
...			
0000079	MVS	International Business Machines Corporation	0000001
...			
...(etc.)			

NOTE: These tables are meant only to be illustrative. In actual implementation, these tables would contain many other fields (columns) than shown here, and the actual structure and contents of the tables would be quite different. For example, in a preferred implementation, there would be a separate vendor table; in a products table, a record would contain only an index pointer into the vendor table.

In one implementation using this approach, a user looking at a license agreement or list of license agreements may also want to look at the inventory/usage information pertaining to a given product. The tool would only have to read the corresponding record number from the license agreement table to have an index pointer into the inventory/usage table.

In another preferred method, a separate table is created consisting of links associating records in a

- 20 -

license information table with corresponding records in a usage information table. Each record of the table consists of two indices, one into one table and one into the other.

5 Figure 5 demonstrates how this association table is applied. One can see that this method provides the additional advantages that it is can be created and maintained without the need to alter either the Inventory/Usage table or the License Agreement table and
10 permits many-to-many associations.

 The inventory/usage repository of an organization's assets (ADB) 18 may contain information which duplicates or closely resembles that contained within the CDB 22. In addition to matching assets in the ADB with those in the
15 CDB, the present invention reconciles such differences.

 Although the CDB and the ADB are conceptually distinct, they can be stored in separate databases or a single database, and that any additional tables or lists that are created by the present invention can be part of
20 that same database, one of the two separate databases, or a third separate database.

 In a preferred embodiment, the present invention provides a facility, the List Associator (LA), for automatically establishing an association between items
25 in the CDB, i.e., agreements, with items in the ADB (i.e., software products, vendors, processors, etc.) whenever product name, product identification number or other identifying feature match.

 In a preferred method, a separate table is created
30 consisting of links associating records in one database

- 21 -

(e.g., the ADB) with corresponding records in another database (e.g., the CDB). An advantage of this method over importing the two types of information into a combined list is that as data contained within the individual repositories is *dynamically updated*, e.g., usage data or the movement of an asset to another physical location, the associations remain always linking the most current data together.

Figure 3 is a computer screen 28 that is an example of product usage viewed by the contractual agreement under which it was acquired.

For example, an agreement for the software product CA-ABEND Aid from Computer Associates Int'l exists in the CDB as record number 0000003, and the product is shown in record number 0000042 of the ADB to be installed on the Development System #2 of Acme Atlanta's Data Processing Center. The LA establishes an association (e.g., 00000030000042) between the two records; thus, if the product is moved to Development System #1, record 0000042 of the ADB is updated. However, since both CDB record 0000003 and association 00000030000042 remain unchanged the user is always presented with the most current information.

The LA creates associations that are many-to-many. Typically, an agreement may include several assets and, an asset may involve multiple agreements (e.g., one for purchase and a second for maintenance). An advantage of this method over importing the two types of information into a combined list is that data fields are not needlessly duplicated.

- 22 -

Optionally, the LA determines possible associations or associations that are most probable. This is based upon the heuristics of evaluating the similarity of item name, identification number, or other such criteria between items in one list and items in another list. Several techniques that can be applied separately or together include: Soundex coding, pattern matching, dictionary (or Knowledge Base) lookup, and/or Fuzzy Logic. The results may be displayed for user selection, confirmation, or rejection. For each item to be associated, the LA provides as many candidates as are appropriate, ranked alphabetically or in likelihood of match. All such candidate associations can be displayed in a distinctive color; the user could tab or scroll directly from one such candidate association to the next and confirm or reject each one in turn, or could highlight and accept or reject several such candidate associations at once.

Optionally, the user can provide a set of rules by which the LA can establish an association between records in separate lists. For example, a rule can be established to treat XYZ Corporation and XYZ International (a wholly-owned subsidiary) as one and the same.

Optionally, the LA provides a facility by which the user specifies an association between one or more items in one list with an item in another list. The user might use a pointing device such as a computer mouse to highlight one or more items on the one list of products and one item on the other list of products. Optionally, the LA would present an opportunity for the user to

- 23 -

confirm the association before it was processed.

Dynamic Detection of Change: Over the course of time, and for a variety of reasons, associations made by the LA may refer to records that have been deleted and, perhaps, reused for other agreements or assets.

Associations for new agreements or assets also need to be made. The user may execute the LA to update these associations and/or this may be scheduled to occur on a periodic basis. However, as a typical installation may have hundreds if not thousands of assets and agreements, this may prove time-consuming and an inconvenience.

Each of the programs that control the ADB and CDB can be modified to notify the LA to update its associations when records are added, deleted, copied or substantially changed (e.g., an asset is disposed of). A disadvantage of this method is that the vendors of these programs may not provide such support.

As an independent notion from the techniques described above, a Dynamic Detection process could directly detect the act of adding, deleting, copying or substantially modifying one or more records in either or both of the ADB and CDB. The Dynamic Detector (DD) is a continuously running process that establishes certain intercepts or hooks into certain operating system or database functions so that when any records in these databases are added, deleted or modified, in any of the several ways in which this could occur, the DD receives control, and is able to analyze the operation. When the DD determines that a record has been added or modified, it may notify the LA to update or create new associations

- 24 -

for those records. If a record has been deleted the LA may decide to perform a more extensive analysis.

Remote Access

The present invention provides various types of remote access capabilities and for a variety of purposes.

A user may provide read-only access to a third party that will use agreement and usage information to negotiate contractual terms with a vendor on behalf of the user.

A user may import information from other sites within the enterprise to view financial and usage data for the purpose of identifying and eliminating unfavorable scenarios (e.g., high-cost contracts with little or no usage), evaluating purchasing and maintenance strategies as well as identifying various cost centers. Similarly, a user may outsource various processing tasks on a sporadic or routine basis and desire to be able to import this usage information.

A user may export information to a service organization that gathers such information from numerous organizations in order to determine industry-wide cost and usage statistics. The user may then evaluate the resulting data in order to determine pricing, purchasing and maintenance strategies.

A user may employ a service bureau to operate the present invention.

In either case, techniques for accessing the appropriate record or records in the one table from a record in the other table are known to those skilled in the art of writing software that accesses data in a

- 25 -

database.

Once these associations are recorded in the database (and the appropriate software is available to make use of these associations, as discussed below), it becomes readily possible for the user to access information automatically and conveniently, as described in more detail in the next section.

Processing, Display, and Output

Initial processing consists of reconciliation and association of the licensing and inventory/usage records as described in the preceding section. Afterwards, the tool provides a number of facilities to process this information for display and output.

These would, among other things, automatically display a list of similar products, while providing the user: the ability to scroll or tab through the complete list to the next such product; an ability to print a list of such products; and the ability to output a flat file with one record for each such product. Information displayed may include:

- agreements for which no inventory information has been associated, as well as a facility for creating a corresponding record in the list of inventoried products, suitably flagged.

- inventoried products for which no agreement information has been associated, as well as a facility for creating a corresponding record in the list of products, suitably flagged, for the agreement portion of the tool.

- products for which both inventory and agreement

- 26 -

information is present.

When an item or list of items providing inventory and/or usage information about a product is on display, the user, via a function key, special keystroke combination, or mouse operation, can invoke a display of the corresponding (associated):

- agreement information record or records.
- usage and/or inventory information record or records.

Where license agreement information is pertinent to inventory and usage information, the tool will optionally calculate and distinctively display such information.

For example, if a license agreement specifies that a product must only be installed on a particular computer, the tool highlights (e.g., in blue) inventory records indicating the product is installed on the authorized computer and highlights (e.g., in red) inventory records indicating the product is installed on an unauthorized computer.

As another example, if a particular agreement licenses a particular product to be installed on computers totaling no more than a certain number of MIPS, the tool totals the MIPS of the processors on which that product has been found (according to the inventory information). If the calculated total is under the licensed capacity, the license is displayed in one way (e.g., in blue), and if the authorized capacity is exceeded, in another way (e.g., in red).

Optionally, the user can select parameters and criteria (such as products on selected processors, at

- 27 -

specified locations, within specified organizational subdivisions, from specified vendors, having usage above or below specified levels over a specified period of time, used by specified users, etc.) whereupon the tool will:

• Distinctively display inventory and usage information according to whether or not user-specified agreement-related criteria are met, such as products on selected processors, at specified locations, within specified organizational subdivisions, from specified vendors, having usage above or below specified levels over a specified period of time, used by specified users, etc. For example, the user can request that the inventory and usage information pertaining to all programs whose maintenance terms are within one month of expiration (according to the corresponding agreement information) be distinctively displayed (e.g., in red), or the display can be filtered to include or exclude all such information.

• Display, for each product (or set of products) found in inventory the relevant license agreement or agreements.

• Display, for each user-selected set of products found in inventory the vendors and/or distributors associated with the agreement or agreements, as well as pertinent information about the vendors and/or distributors, such as their Internet Web site addresses (optionally, with the facility to open a browser Window displaying information from the appropriate Web site).

While the user can choose to select a subset of all

- 28 -

products, he or she can choose to include all products.

- Display, for each user-selected set of products found in inventory contact information (such as name, address, telephone numbers, email addresses) pertaining to the agreement or agreements.

- Display, for each user-selected set of products found in inventory important dates associated with the agreement or agreements (such as cancellation and renewal dates).

- Display, for each user-selected set of products found in inventory invoice and/or payment forecast, budget, and/or history information.

- Display, for all user-selected set of products found in inventory the products that either do or do not have an agreement associated with them.

- Display for all agreements, all products that have not been found in inventory.

- Display, for all products whose maintenance contracts are within a user-specified time of cancellation or renewal, the recent and current usage levels.

- Display for each agreement on which processors, if any, the product is installed and how much it has been used. Included in the display of usage information is which users, on which processors, at what locations, by what organizational subdivisions, when it was used (by date, by hours of use, days of the week of use, etc.), etc. Optionally, all such information can be displayed for each version and/or release of the product identified.

- 29 -

- Display for each agreement, usage trends and projections, including by costs incurred over specified time periods.

- Display for each agreement that specifies a
5 maximum number of MIPS (a measuring of relative processing power), the processors on which the product is installed.

- Display for each agreement that specifies a particular number of processors the number of processors
10 on which it is installed, and whether that number is higher or lower than the agreement authorizes.

- Display for each agreement that specifies a particular number of processors the number of processors on which it has been used during a user-specified period
15 of time, and whether that number is higher or lower than the agreement authorizes.

- Display for each agreement that specifies that the product can be installed on one or more particular processors, the processors is the product installed, and
20 whether any such use is not authorized by the agreement.

- Display for each agreement that specifies that the product can be installed on one or more particular processors for which the user has an agreement, the processors on which the product has been used during a
25 user-specified period of time, and whether any such use is not authorized by the agreement.

- Display, for each agreement that authorizes particular (or multiple) versions and/or releases of the product the versions and/or releases installed,
30 indicating any of these are not authorized by the

- 30 -

agreement.

- Display charts and graphs of any or all of the user-selected and user-specified information selected for display, with records meeting specified criteria

5 distinctively highlighted.

Information presented to the user via interactive display can also be used to

- Produce printed reports of user-selected and user-specified information, distinctively highlighting

10 records that meet the specified criteria.

- Export records for use by other programs and products of user-selected and user-specified information.

Furthermore, the invention can, among other things:

- Display, for each user-selected set of products

15 found in inventory invoice and/or payment forecast, budget, and/or history information.

- Display for each agreement on which processors, if any, the product is installed and how much it has been used. Included in the display of usage information is

20 which users, on which processors, at what locations, by what organizational subdivisions, when it was used (by date, by hours of use days of the week of use, etc.), etc. Optionally, all such information can be displayed for each version and/or release of the product

25 identified.

- Display for each agreement, usage trends and projections.

Since costing information is included in the License Agreement Table, further aspects relating to cost are

30 possible with this invention. Such features include:

- 31 -

- Display, for all products whose maintenance contracts are within a user-specified time of cancellation or renewal, the anticipated costs based upon recent and current usage levels.

5 • Display for each agreement on which processors, if any, the product is installed and the actual licensing costs based upon how much it has been used over selected time periods, e.g. prior three, or six etc. months.

10 Included in the display of this information is which users, on which processors, at what locations, by what organizational subdivisions, when it was used (by date, by hours of use, days of the weeks of use, etc.), etc. Optionally, all such information can be displayed for each version and/or release of the product identified.

15 • Display for each agreement, licensing cost trends and projections.

20 Certain output aspects of the present invention are illustrated in Figure 6 which shows the usage of selected software assets organized according to agreement and location.

25 When a software tool, e.g. the data acquisition tool, is said to be "independent," it is to be understood that the tool stands alone, operating independently of the other tools, for example, as when using the present assignee's SoftAudit® software to gather inventory and usage data. Relative to data bases, the term means that one data base can be modified without being concerned about the contents of other data bases, without destroying the usefulness or operability of the system of

30 the present invention.

- 32 -

A conceptual data model of certain aspects of the present invention is illustrated in Figure 4 which shows various correlations of information. In the figure, a straight line denotes a one-to-one database relationship.

5 A line that is dotted on both ends denotes a many-to-many relationship. A line that is only singly dotted denotes a many-to-one relationship.

10 Although the present invention has been described in relation to particular embodiments thereof, many other variations and modifications and other uses will become apparent to those skilled in the art. It is preferred, therefore, that the present invention be limited not by the specific disclosure herein, but only by the appended claims.

- 33 -

WHAT IS CLAIMED IS:

1. A knowledge base facility which correlates software product data relating to software products installed on a computer with related software agreement data, the knowledge base facility comprising:

5 an independent first database holding a plurality of software product data records;

an independent second database holding a plurality of software agreement data records;

10 linking data which links each of a plurality of said software product data records with one or more corresponding records of said software agreement data records and vice versa, said linking data permitting a user to identify associated data records of said software product data records and said software agreement data records;

15 a query tool that receives and acts on queries from a user for linked data records; and

20 a data outputting facility that outputs records from said software product data records and said software agreement data records in a response to said query tool.

2. The knowledge base facility of claim 1, in which the software product data records comprise both inventory and/or usage data records.

3. The knowledge base facility of claim 1, in which the software agreement data records comprise records at least some of which are organized as a plurality of fields selected from a group containing:

- 34 -

5 time and date of transaction; name, address, phone
number, e-mail address of vendor; description of asset;
manufacturer, model, year, options; warranties, contract
terms and conditions; and contract information; licensing
fees, number of persons authorized; MIPS authorized;
10 number of invocations authorized; annual maintenance
fees; and product/sub-product identifications.

4. The knowledge base facility of claim 3, in
which the linking data includes links to data files
containing substantially the entireties of software
contract agreements that are associated with one or more
5 of the software product data records.

5. The knowledge base facility of claim 1,
including a link data updater which maintains the linking
data upon dynamic updating of data records in one or the
other or both of said first and second databases.

6. The knowledge base facility of claim 1, in
which the linking data contains one-to-many and many-to-
many data links.

7. The knowledge base facility of claim 1, in
which the linking data is created by one or more of:
Soundex Coding, pattern matching, dictionary look-up
and/or Fuzzy Logic.

8. A method for correlating software product data
relating to software products installed on a computer

- 35 -

with related software agreement data, the method comprising the steps of:

- 5 receiving with a first software facility the software product data;
- receiving with a second software facility the software agreement data;
- 10 correlating substantially automatically with a third software facility the software product data with the software agreement data; and
- outputting results of said correlation with a fourth software facility.

9. The method of claim 8, including performing the method substantially automatically.

10. The method of claim 9, in which the software product data comprises inventory and/or usage data.

11. The method of claim 9, in which the software agreement data comprises a table containing summary information concerning agreements aggregated as various data fields, and the fields are selected from the group
- 5 containing: time and date of transaction; name, address, phone number, e-mail address of vendor; description of asset; manufacturer, model, year, options; warranties, contract terms and conditions; contract information, licensing fees, number of persons authorized, MIPS
- 10 authorized, number of invocations authorized, annual maintenance fees, product/sub-product identification.

- 36 -

12. The method of claim 9, in which the correlating step is carried out so that at least 25% of one of the software product data and software agreement data has been correlated.

13. The method of claim 10, in which the third software facility correlates all software products meeting user-specified criteria.

14. The method of claim 12, including selecting for processing by the correlating software, products with usage above or below a user specified level over a user-specified time period.

15. The method of claim 12, in which the user-specified criteria includes software products whose maintenance contracts are within a user-specified time of cancellation or renewal.

16. The method of claim 12, in which the user-specified criteria includes (a) products for which the user has a license agreement and/or (b) a processor on which products are installed and/or (c) at least one
5 location of software products.

17. A substantially automatic system for correlating software product data concerning software products installed on a computer with related software agreement data, the system comprising:

5 a first software facility for receiving the software

product data;

a second software facility for receiving the software agreement data;

10 a third software facility for correlating substantially automatically the software product data and the software agreement data with one another; and

a fourth software facility for outputting results of said correlation.

18. The invention of any one of claims 1, 8 and 17, in which the computer is a mainframe computer or a network of computers.

19. The system of claim 17, in which the agreement data contains a plurality of data fields selected from a group including:

5 time and date of transaction; name, address, phone number, e-mail address of vendor; description of asset; manufacturer, model, year, options; warranties, contract terms and conditions; contract information, fee structure; number of concurrent users; renewal dates; MIPS authorization; authorized CPUs; product name; and
10 discounts.

20. The system of claim 17, in which the third software facility comprises a reconciliation and association facility which includes a plurality of software tools.

21. The system of claim 20, in which the

- 38 -

reconciliation and association tools include a tool for resolving inconsistencies and mismatches in product names between the software product data and the software agreement data.

22. The system of claim 17, in which the third software facility is effective for correlating at least 25% of the data contained in one of the software product data and the software agreement product data.

23. The invention of any one of claims 1 or 8, including a link data updater that is enabled to run on a periodic basis, or continuously or in response to alternations to the inventory of software products.

24. The method of claim 9, including maintaining the software agreement data on a second computer located remotely from the computer on which the software products are installed.

25. The method of claim 24, in which said software agreement data is representative of industry-wide licensing cost statistics, for being compared to specific software agreement data applicable to the software products installed on the computer.

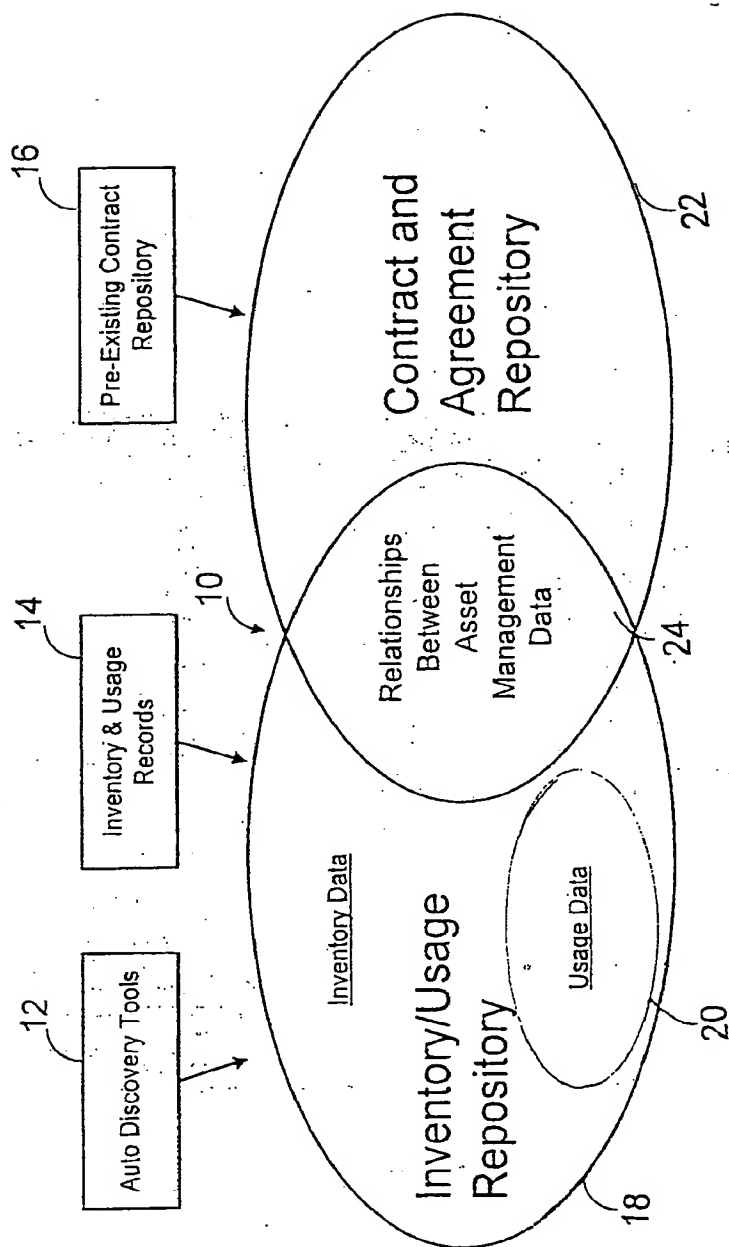
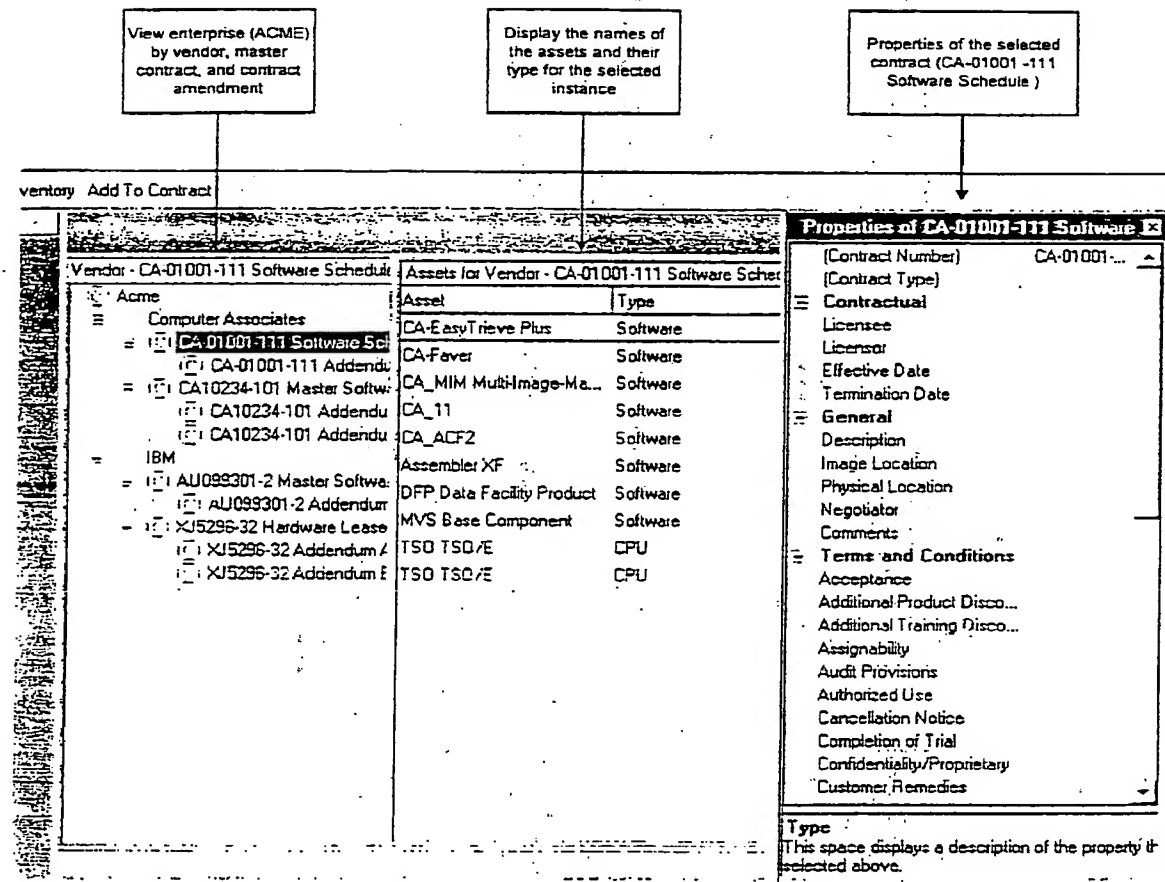


Figure 1



26

Figure 2

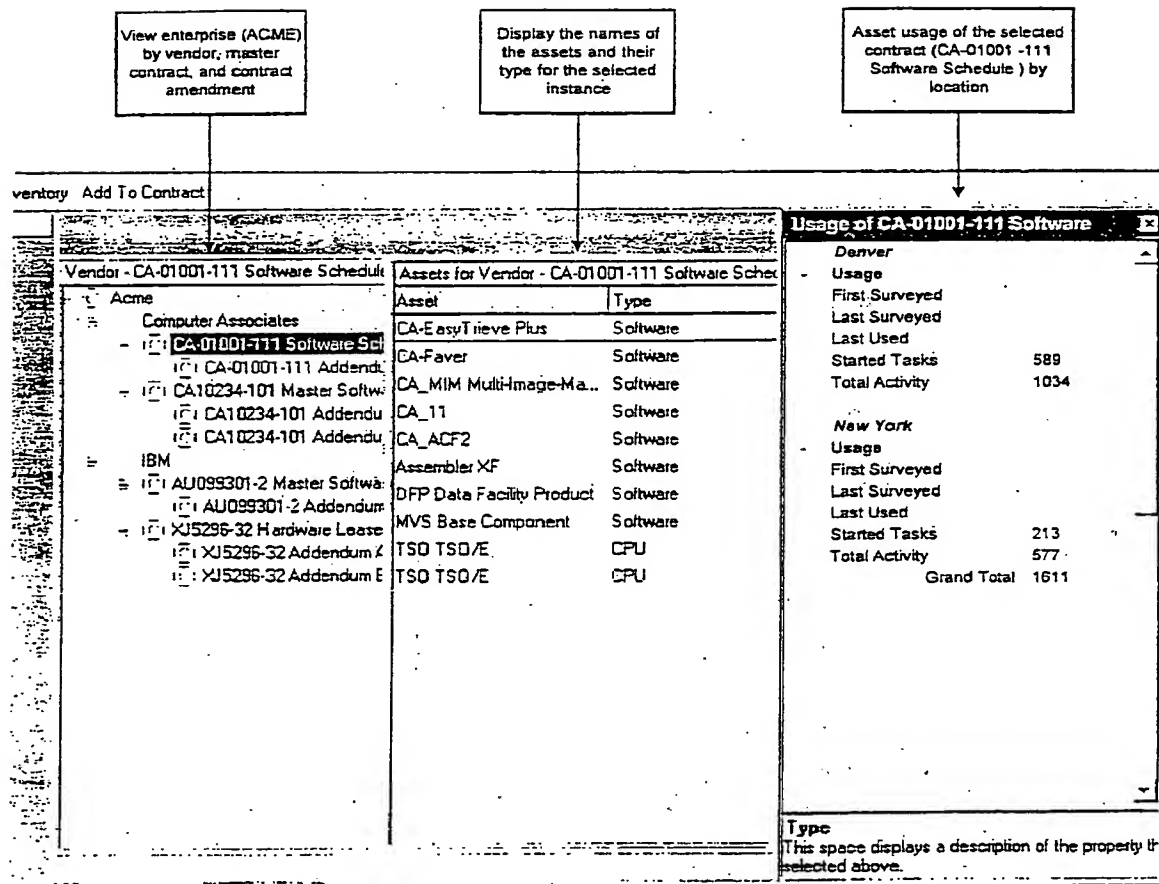


Figure 3

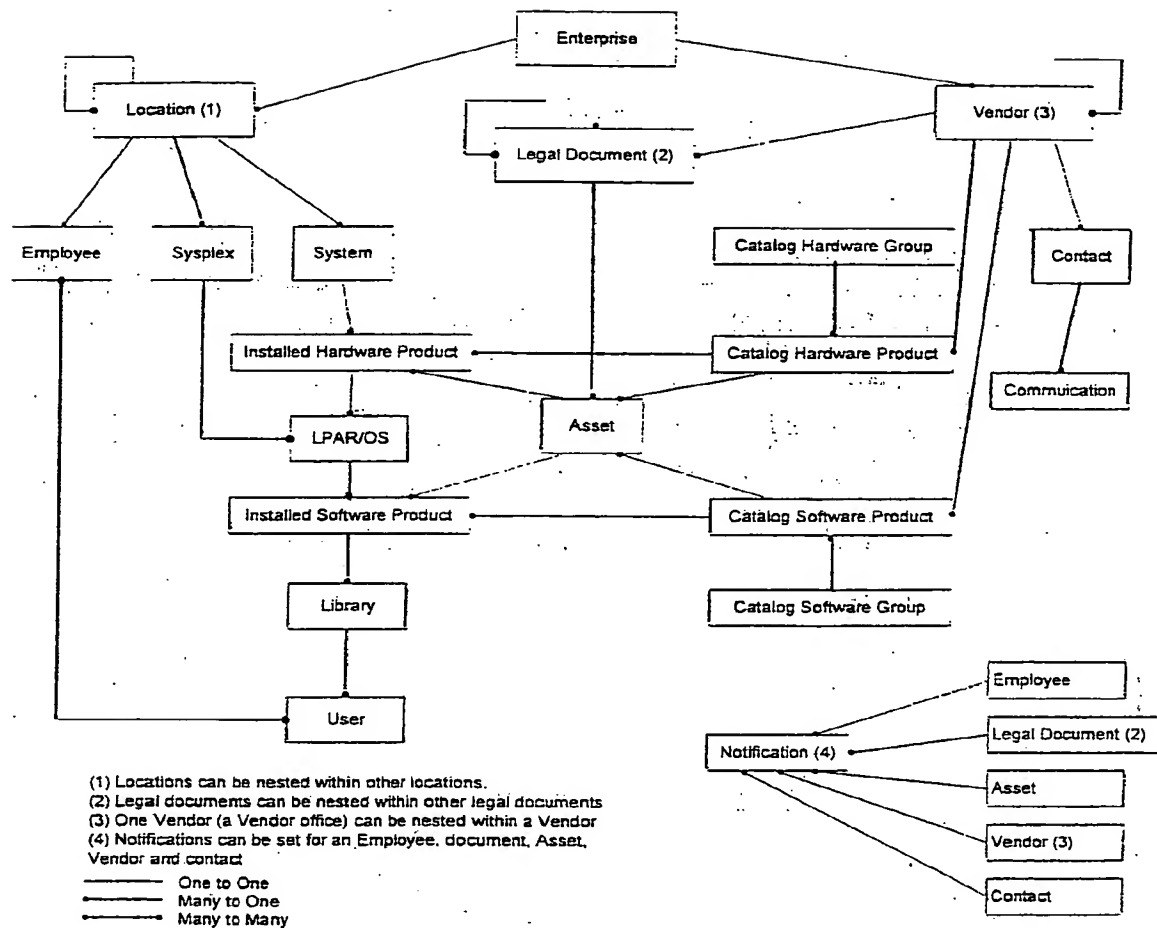


Figure 4

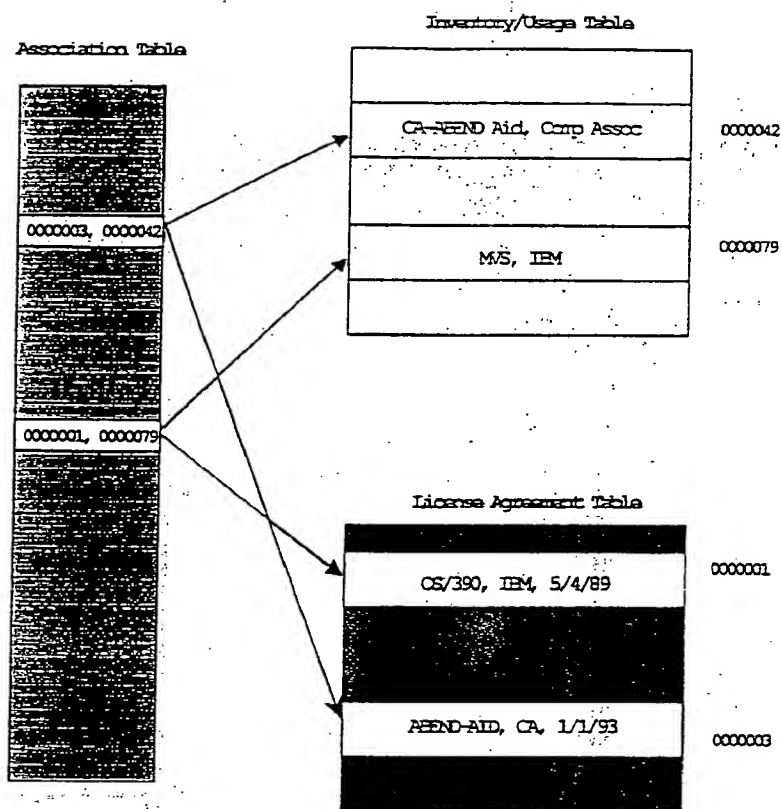


Figure 5

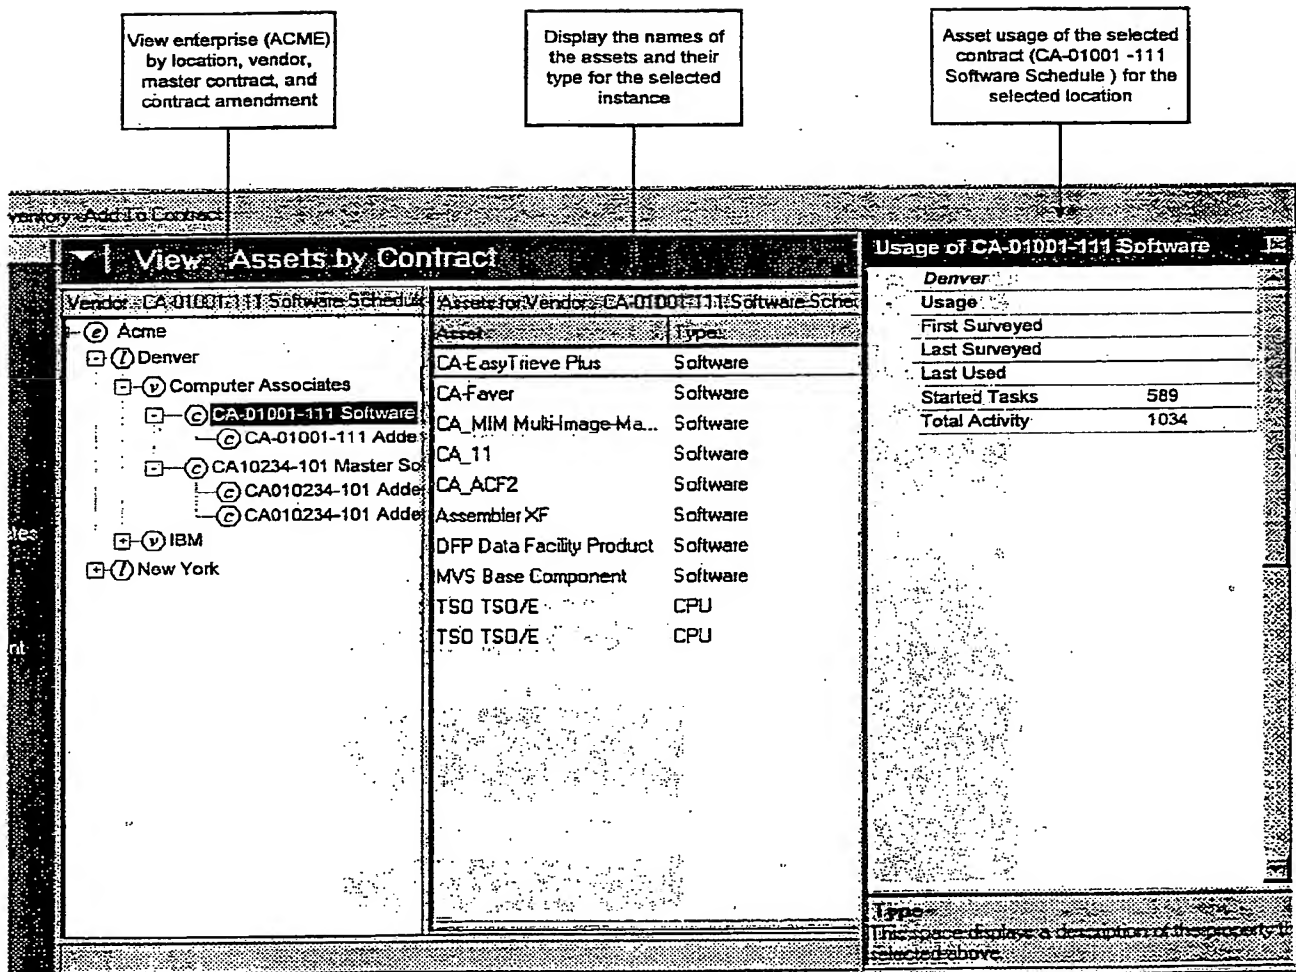


Figure 6 - Asset Usage by Location and Contract

INTERNATIONAL SEARCH REPORT

Intel onal Application No
PCT/US 00/05664

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 G06F1/00

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, PAJ

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 0 852 349 A (ISOGON CORP) 8 July 1998 (1998-07-08) abstract; figures 1-4 column 1, line 29 - line 52 column 3, line 20 -column 12, line 39 column 14, line 46 -column 15, line 17 column 18, line 24 - line 39 claims 31,32	8-13, 16-19, 22,24
Y		1-6,23
Y	WO 97 37316 A (COPYRIGHT CLEARANCE CENTER INC ;JOHNSON WOODROW W (US); ATKINS CHR) 9 October 1997 (1997-10-09) abstract; figures 4-7 page 3, line 25 -page 5, line 6 page 10, line 6 -page 16, line 4 ----- -/-	1-6,23

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
- "&" document member of the same patent family

Date of the actual completion of the international search

5 July 2000

Date of mailing of the international search report

12/07/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Sigolo, A

INTERNATIONAL SEARCH REPORT

Inter. Patent Application No

PCT/US 00/05664

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5 202 977 A (PASETES JR EMMANUEL K ET AL) 13 April 1993 (1993-04-13) column 17, line 31 - line 40 -----	7,20

INTERNATIONAL SEARCH REPORT

information on patent family members

International Application No

PCT/US 00/05664

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0852349 A	08-07-1998	US 6029145 A	22-02-2000
WO 9737316 A	09-10-1997	US 5991876 A	23-11-1999
		AU 719229 B	04-05-2000
		AU 2425297 A	22-10-1997
		CA 2235276 A	09-10-1997
		EP 0829059 A	18-03-1998
US 5202977 A	13-04-1993	AT 161342 T	15-01-1998
		AU 8314691 A	04-02-1992
		DE 69128483 D	29-01-1998
		EP 0553098 A	04-08-1993
		WO 9201251 A	23-01-1992